

METHODS AND SYSTEMS FOR PROCESSING INSTANT MESSENGER MESSAGES

FIELD OF THE INVENTION

[0001] The invention generally relates to search engines. More particularly, the invention relates to methods and systems for processing instant messenger messages.

BACKGROUND OF THE INVENTION

[0002] Users generate and access a large number of articles, such as emails, web pages, word processing documents, spreadsheet documents, instant messenger messages, and presentation documents, using a client device, such as a personal computer, personal digital assistant, or mobile phone. Some articles are stored on one or more storage devices coupled to, accessible by, or otherwise associated with the client device(s). Users sometimes wish to search the storage device(s) for articles.

[0003] Conventional client-device search applications may significantly degrade the performance of the client device. For example, certain conventional client-device search applications typically use batch processing to index all articles, which can result in noticeably slower performance of the client device during the batch processing. Additionally, batch processing occurs only periodically. Therefore, when a user performs a search, the most recent articles are sometimes not included in the results. Moreover, if the batch processing is scheduled for a time when the client device is not operational and

is thus not performed for an extended period of time, the index of articles associated with the client device can become outdated. Conventional client-device search applications can also need to rebuild the index at each batch processing or build new partial indexes and perform a merge operation that can use a lot of client-device resources.

Conventional client-device search applications also sometimes use a great deal of system resources when operational, resulting in slower performance of the client device.

[0004] Generally, conventional search engines do not index instant messenger messages and/or allow a user to search instant messenger messages. Some applications exist that store and/or log a user's instant messages, but these applications generally do not index the user's instant messenger messages.

[0005] Additionally, conventional client-device search applications can require an explicit search query from a user to generate results, and may be limited to file names or the contents of a particular application's files.

SUMMARY

[0006] Embodiments of the present invention comprise methods and systems for processing instant messenger messages. In one embodiment, an instant messenger event is captured by compiling event data associated with at least one instant messenger message, the instant messenger event is associated with a conversation, and at least some of the event data associated with the instant messenger event is indexed. In one embodiment, a search query is received, and the conversation is identified as relevant to

the search query. In one embodiment, capturing the instant message event can comprise identifying instant messenger activity associated with an instant messenger application on a client device, identifying the instant messenger event, and compiling the instant messenger event from at least some of the event data.

[0007] These exemplary embodiments are mentioned not to limit or define the invention, but to provide examples of embodiments of the invention to aid understanding thereof. Exemplary embodiments are discussed in the Detailed Description, and further description of the invention is provided there. Advantages offered by the various embodiments of the present invention may be further understood by examining this specification.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] These and other features, aspects, and advantages of the present invention are better understood when the following Detailed Description is read with reference to the accompanying drawings, wherein:

[0009] Figure 1 is a diagram illustrating an exemplary environment in which one embodiment of the present invention may operate;

[0010] Figure 2 is a flow diagram illustrating an exemplary method of capturing and processing event data associated with a client device in one embodiment of the present invention; and

[0011] Figure 3 is a flow diagram illustrating an exemplary method of capturing an instant messenger event in one embodiment of the present invention; and

[0012] Figure 4 is a flow diagram illustrating an exemplary method of indexing an instant messenger event in one embodiment of the present invention.

DETAILED DESCRIPTION OF SPECIFIC EMBODIMENTS

[0013] Referring now to the drawings in which like numerals indicate like elements throughout the several figures, Figure 1 is a block diagram illustrating an exemplary environment for implementation of an embodiment of the present invention. While the environment shown in Figure 1 reflects a client-side search engine architecture embodiment, other embodiments are possible. The system 100 shown in Figure 1 includes multiple client devices 102a-n that can communicate with a server device 150 over a network 106. The network 106 shown in Figure 1 comprises the Internet. In other embodiments, other networks, such as an intranet, may be used instead. Moreover, methods according to the present invention may operate within a single client device that does not communicate with a server device or a network.

[0014] The client devices 102a-n shown in Figure 1 each include a computer-readable medium 108. The embodiment shown in Figure 1 includes a random access memory (RAM) 108 coupled to a processor 110. The processor 110 executes computer-executable program instructions stored in memory 108. Such processors may include a microprocessor, an ASIC, state machines, or other processor, and can be any of a number

of suitable computer processors, such as processors from Intel Corporation of Santa Clara, California and Motorola Corporation of Schaumburg, Illinois. Such processors include, or may be in communication with, media, for example computer-readable media, which stores instructions that, when executed by the processor, cause the processor to perform the steps described herein. Embodiments of computer-readable media include, but are not limited to, an electronic, optical, magnetic, or other storage or transmission device capable of providing a processor, such as the processor 110 of client 102a, with computer-readable instructions. Other examples of suitable media include, but are not limited to, a floppy disk, CD-ROM, DVD, magnetic disk, memory chip, ROM, RAM, an ASIC, a configured processor, all optical media, all magnetic tape or other magnetic media, or any other medium from which a computer processor can read instructions. Also, various other forms of computer-readable media may transmit or carry instructions to a computer, including a router, private or public network, or other transmission device or channel, both wired and wireless. The instructions may comprise code from any suitable computer-programming language, including, for example, C, C++, C#, Visual Basic, Java, Python, Perl, and JavaScript.

[0015] Client devices 102a-n can be coupled to a network 106, or alternatively, can be stand alone machines. Client devices 102a-n may also include a number of external or internal devices such as a mouse, a CD-ROM, DVD, a keyboard, a display device, or other input or output devices. Examples of client devices 102a-n are personal computers, digital assistants, personal digital assistants, cellular phones, mobile phones, smart

phones, pagers, digital tablets, laptop computers, Internet appliances, and other processor-based devices. In general, the client devices 102a-n may be any type of processor-based platform that operates on any suitable operating system, such as Microsoft® Windows® or Linux, capable of supporting one or more client application programs. For example, the client device 102a can comprise a personal computer executing client application programs, also known as client applications 120. The client applications 120 can be contained in memory 108 and can include, for example, a word processing application, a spreadsheet application, an email application, an instant messenger application, a presentation application, an Internet browser application, a calendar/organizer application, a video playing application, an audio playing application, an image display application, a file management program, an operating system shell, and other applications capable of being executed by a client device. Client applications may also include client-side applications that interact with or accesses other applications (such as, for example, a web-browser executing on the client device 102a that interacts with a remote e-mail server to access e-mail).

[0016] The user 112a can interact with the various client applications 120 and articles associated with the client applications 120 via various input and output devices of the client device 102a. Articles include, for example, word processor documents, spreadsheet documents, presentation documents, emails, instant messenger messages, database entries, calendar entries, appointment entries, task manager entries, source code files, and other client application program content, files, messages, items, web pages of

various formats, such as HTML, XML, XHTML, Portable Document Format (PDF) files, and media files, such as image files, audio files, and video files, or any other documents or items or groups of documents or items or information of any suitable type whatsoever.

[0017] The user's 112a interaction with articles, the client applications 120, and the client device 102a creates event data that may be observed, recorded, analyzed or otherwise used. An event can be any occurrence possible associated with an article, client application 120, or client device 102a, such as inputting text in an article, displaying an article on a display device, sending an article, receiving an article, manipulating an input device, opening an article, saving an article, printing an article, closing an article, opening a client application program, closing a client application program, idle time, processor load, disk access, memory usage, bringing a client application program to the foreground, changing visual display details of the application (such as resizing or minimizing) and any other suitable occurrence associated with an article, a client application program, or the client device whatsoever. Additionally, event data can be generated when the client device 102a interacts with an article independent of the user 112a, such as when receiving an email or performing a scheduled task.

[0018] The memory 108 of the client device 102a can also contain a capture processor 124, a queue 126, and a search engine 122. The client device 102a can also contain or is in communication with a data store 140. The capture processor 124 can capture events and pass them to the queue 126. The queue 126 can pass the captured events to the search engine 122 or the search engine 122 can retrieve new events from the

queue 126. In one embodiment, the queue 126 notifies the search engine 122 when a new event arrives in the queue 126 and the search engine 122 retrieves the event (or events) from the queue 126 when the search engine 122 is ready to process the event (or events). When the search engine receives an event it can be processed and can be stored in the data store 140. The search engine 122 can receive an explicit query from the user 112a or generate an implicit query and it can retrieve information from the data store 140 in response to the query. In another embodiment, the queue is located in the search engine 122. In still another embodiment, the client device 102a does not have a queue and the events are passed from the capture processor 124 directly to the search engine 122. According to other embodiments, the event data is transferred using an information exchange protocol. The information exchange protocol can comprise, for example, any suitable rule or convention facilitating data exchange, and can include, for example, any one of the following communication mechanisms: Extensible Markup Language – Remote Procedure Calling protocol (XML/RPC), Hypertext Transfer Protocol (HTTP), Simple Object Access Protocol (SOAP), shared memory, sockets, local or remote procedure calling, or any other suitable information exchange mechanism.

[0019] The capture processor 124 can capture an event by identifying and compiling event data associated with an event. Examples of events include sending or receiving an instant messenger message, a user viewing a web page, saving a word processing document, printing a spreadsheet document, inputting text to compose or edit an email, opening a presentation application, closing an instant messenger application, entering a

keystroke, moving the mouse, and hovering the mouse over a hyperlink. An example of event data captured by the capture processor 124 for an event involving the receipt of an instant messenger message by the user 112a can comprise the sender of the message, the recipients of the message, the time and date the message was received, the content of the message and a conversation ID. A conversation ID can be used to associate messages that form a conversation and can be provided by the instant messenger application or can be generated by the capture processor 124. A conversation can be one or more messages between the user 112a and at least one other user until the user 112a logs out of or closes the instant messenger application or the instant messenger application is inactive for a certain period of time (for example, 30 minutes).

[0020] In the embodiment shown in Figure 1, the capture processor 124 comprises multiple capture components. For example, the capture processor 124 shown in Figure 1 comprises a separate capture component for each client application in order to capture events associated with each application. The capture processor 124 can also comprises a separate capture component that monitors overall network activity in order to capture event data associated with network activity, such as the receipt or sending of an instant messenger message. The capture processor 124 shown in Figure 1 also can comprise a separate client device capture component that monitors overall client device performance data, such as processor load, idle time, disk access, the client applications in use, and the amount of memory available. The capture processor 124 shown in Figure 1 also comprises a separate capture component to monitor and capture keystrokes input by the

user and a separate capture component to monitor and capture items, such as text, displayed on a display device associated with the client device 102a. An individual capture component can monitor multiple client applications and multiple capture components can monitor different aspects of a single client application.

[0021] In one embodiment, the capture processor 124, through the individual capture components, can monitor activity on the client device and can capture events by a generalized event definition and registration mechanism, such as an event schema. Each capture component can define its own event schema or can use a predefined one. Event schemas can differ depending on the client application or activity the capture component is monitoring. Generally, the event schema can describe the format for an event, for example, by providing fields for event data associated with the event (such as the time of the event) and fields related to any associated article (such as the title) as well as the content of any associated article (such as the document body). An event schema can describe the format for any suitable event data that relates to an event. For example, an event schema for an instant messenger message event sent by the user 112a can include a recipient or list of recipients, the time sent, the date sent, content of the message, and a conversation ID. An event schema for a web page currently being viewed by a user can include the Uniform Resource Locator or URL of the web page, the time being viewed, and the content of the web page. An event schema for a word processing document being saved by a user can include the title of the document, the time saved, the location of the document, the format of the document, the text of the document, and a pointer to the

location of the document. More generally, an event schema can describe the state of the system around the time of the event. For example, an event schema can contain a URL for a web page event associated with a previous web page that the user navigated from. In addition, event schema can describe fields with more complicated structure like lists. For example, an event schema can contain fields that list multiple recipients. An event schema can also contain optional fields so that an application can include additional event data if desired.

[0022] The capture processor 124 can capture events occurring presently (or “real-time events”) and can capture events that have occurred in the past (or “historical events”). Real-time events can be “indexable” or “non-indexable”. In one embodiment, the search engine 122 indexes indexable real-time events, but does not index non-indexable real-time events. The search engine 122 may determine whether to index an event based on the importance of the event. Indexable real-time events can be more important events associated with an article, such as viewing a web page, loading or saving a file, and receiving or sending an instant message or email. Non-indexable events can be deemed not important enough by the search engine 122 to index and store the event, such as moving the mouse or selecting a portion of text in an article. Non-indexable events can be used by the search engine 122 to update the current user state. While all real-time events can relate to what the user is currently doing (or the current user state), indexable real-time events can be indexed and stored in the data store 140. Alternatively, the search engine 122 can index all real-time events. Real-time events can

include, for example, sending or receiving an article, such as an instant messenger message, examining a portion of an article, such as selecting a portion of text or moving a mouse over a portion of a web page, changing an article, such as typing a word in an email or pasting a sentence in a word processing document, closing an article, such as closing an instant messenger display area or changing an email message being viewed, loading, saving, opening, or viewing an article, such as a word processing document, web page, or email, listening to or saving an MP3 file or other audio/video file, or updating the metadata of an article, such as book marking a web page, printing a presentation document, deleting a word processing document, or moving a spreadsheet document.

[0023] Historical events are similar to indexable real-time events except that the event occurred before the installation of the search engine 122 or was otherwise not captured, because, for example, the search engine 122 was not operational for a period of time while the client device 102a was operational or because no capture component existed for a specific type of historical event at the time the event took place. Examples of historical events include the user's saved word processing documents, media files, presentation documents, calendar entries, and spreadsheet documents, the emails in a user's inbox, and the web pages book marked by the user. The capture processor 124 can capture historical events by periodically crawling the memory 108 and any associated data storage device for events not previously captured by the capture processor 124. The capture processor 124 can also capture historical events by requesting certain client applications, such as a web browser or an email application, to retrieve articles and other

associated information. For example, the capture processor 124 can request that the web browser application obtain all viewed web pages by the user or request that the email application obtain all email messages associated with the user. These articles may not currently exist in memory 108 or on a storage device of the client device 102a. For example, the email application may have to retrieve emails from a server device. In one embodiment, the search engine 122 indexes historical events.

[0024] In the embodiment shown in Figure 1, events captured by the capture processor 124 are sent to the queue 126 in the format described by an event schema. The capture processor 124 can also send performance data to the queue 126. Examples of performance data include current processor load, average processor load over a predetermined period of time, idle time, disk access, the client applications in use, and the amount of memory available. Performance data can also be provided by specific performance monitoring components, some of which may be part of the search engine 122, for example. The performance data in the queue 126 can be retrieved by the search engine 122 and the capture components of the capture processor 124. For example, capture components can retrieve the performance data to alter how many events are sent to the queue 126 or how detailed the events are that are sent (fewer or smaller events when the system is busy) or how frequently events are sent (events are sent less often when the system is busy or there are too many events waiting to be processed). The search engine 122 can use performance data to determine when it indexes various events and when and how often it issues implicit queries.

[0025] In one embodiment, the queue 126 holds events until the search engine 122 is ready to process an event or events. Alternatively, the queue 126 uses the performance data to help determine how quickly to provide the events to the search engine 122. The queue 126 can comprise one or more separate queues including a user state queue and an index queue. The index queue can queue indexable events, for example. Alternatively, the queue 126 can have additional queues or comprise a single queue. The queue 126 can be implemented as a circular priority queue using memory mapped files. The queue can be a multiple-priority queue where higher priority events are served before lower priority events, and other components may be able to specify the type of events they are interested in. Generally, real-time events can be given higher priority than historical events, and indexable events can be given higher priority than non-indexable real-time events. Other implementations of the queue 126 are possible. In another embodiment, the client device 102a does not have a queue 126. In this embodiment, events are passed directly from the capture processor 124 to the search engine 122. In other embodiments, events can be transferred between the capture components and the search engine using suitable information exchange mechanisms such as: Extensible Markup Language – Remote Procedure Calling protocol (XML/RPC), Hypertext Transfer Protocol (HTTP), Simple Object Access Protocol (SOAP), shared memory, sockets, local or remote procedure calling, or any other suitable information exchange mechanism.

[0026] The search engine 122 can contain an indexer 130, a query system 132, and a formatter 134. The query system 132 can retrieve all real-time events and performance

data from the queue 126. The query system 132 can use performance data and real-time events to update the current user state and generate an implicit query. An implicit query can be an automatically generated query based on the current user state. The query system 132 can also receive and process explicit queries from the user 112a.

Performance data can also be retrieved by the search engine 122 from the queue 126 for use in determining the amount of activity possible by the search engine 122.

[0027] In the embodiment shown in Figure 1, indexable real-time events and historical events (indexable events) are retrieved from the queue 126 by the indexer 130. Alternatively, the queue 126 may send the indexable events to the indexer 130. The indexer 130 can index the indexable events and can send them to the data store 140 where they are stored. The data store 140 can be any type of computer-readable media and can be integrated with the client device 102a, such as a hard drive, or external to the client device 102a, such as an external hard drive or on another data storage device accessed through the network 106. The data store can be one or more logical or physical storage areas. In one embodiment, the data store 140 can be in memory 108. The data store 140 may facilitate one or combination of methods for storing data, including without limitation, arrays, hash tables, lists, and pairs, and may include compression and encryption. In the embodiment shown in Figure 1, the data store comprises an index 142, a database 144 and a repository 146.

[0028] In the embodiment shown in Figure 1, when the indexer 130 receives an event, the indexer 130 can determine, from the event schema, terms (if any) associated

with the event, the time of the event (if available), images (if any) associated with the event, and/or other information defining the event. The indexer 130 can also determine if the event relates to other events and associate the event with related events. For example, for a received instant messenger message event, the indexer can associate the message event with other message events from the same conversation. The messages from the same conversation can be associated with each other in a conversation object, which can be stored in the data store 140.

[0029] The indexer 130 can send and incorporate the terms and times, associated with the event in the index 142 of the data store 140. The event can be sent to the database 144 for storage and the content of the associated article and any associated images can be stored in the repository 146. The conversation object associated with instant messenger messages can be stored in the database 144.

[0030] In the embodiment shown in Figure 1, a user 112a can input an explicit query into a search engine interface displayed on the client device 102a, which is received by the search engine 122. The search engine 122 can also generate an implicit query based on a current user state, which can be determined by the query system 132 from real-time events. Based on the query, the query system 132 can locate relevant information in the data store 140 and provide a result set. In one embodiment, the result set comprises article identifiers for articles associated with the client applications 120 or client articles. Client articles include articles associated with the user 112a or client device 102a, such as the user's emails, word processing documents, instant messenger messages, previously

viewed web pages and any other article or portion of an article associated with the client device 102a or user 112a. An article identifier may be, for example, a Uniform Resource Locator (URL), a file name, a link, an icon, a path for a local file, or other suitable information that may identify an article. In another embodiment, the result set also comprises article identifiers for articles located on the network 106 or network articles located by a search engine on a server device. Network articles include articles located on the network 106 not previously viewed or otherwise referenced by the user 112a, such as web pages not previously viewed by the user 112a.

[0031] The formatter 134 can receive the search result set from the query system 132 of the search engine 122 and can format the results for output to a display processor 128. In one embodiment, the formatter 134 can format the results in XML, HTML, or tab delineated text. The display processor 128 can be contained in memory 108 and can control the display of the result set on a display device associated with the client device 102a. The display processor 128 may comprise various components. For example, in one embodiment, the display processor 128 comprises a Hypertext Transfer Protocol (HTTP) server that receives requests for information and responds by constructing and transmitting Hypertext Markup Language (HTML) pages. In one such embodiment, the HTTP server comprises a scaled-down version of the Apache Web server. The display processor 128 can be associated with a set of APIs to allow various applications to receive the results and display them in various formats. The display APIs can be

implemented in various ways, including as, for example, DLL exports, COM interface, VB, JAVA, or .NET libraries, or as a web service.

[0032] Through the client devices 102a-n, users 112a-n can communicate over the network 106, with each other and with other systems and devices coupled to the network 106. As shown in Figure 1, a server device 150 can be coupled to the network 106. In the embodiment shown in Figure 1, the search engine 122 can transmit a search query comprised of an explicit or implicit query or both to the server device 150. The user 112a can also enter a search query in a search engine interface, which can be transmitted to the server device 150 by the client device 102a via the network 106. In another embodiment, the query signal may instead be sent to a proxy server (not shown), which then transmits the query signal to server device 150. Other configurations are also possible.

[0033] The server device 150 can include a server executing a search engine application program, such as the Google™ search engine. In other embodiments, the server device 150 can comprise a related information server or an advertising server. Similar to the client devices 102a-n, the server device 150 can include a processor 160 coupled to a computer-readable memory 162. Server device 150, depicted as a single computer system, may be implemented as a network of computer processors. Examples of a server device 150 are servers, mainframe computers, networked computers, a processor-based device, and similar types of systems and devices. The server processor 160 can be any of a number of computer processors, such as processors from Intel Corporation of Santa Clara, California and Motorola Corporation of Schaumburg,

Illinois. In another embodiment, the server device 150 may exist on a client-device. In still another embodiment, there can be multiple server devices 150.

[0034] Memory 162 contains the search engine application program, also known as a network search engine 170. The search engine 170 can locate relevant information from the network 106 in response to a search query from a client device 102a. The search engine 170 then can provide a result set to the client device 102a via the network 106. The result set can comprise one or more article identifiers. An article identifier may be, for example, a Uniform Resource Locator (URL), a file name, a link, an icon, a path for a local file, or anything else that identifies an article. In one embodiment, an article identifier can comprise a URL associated with an article.

[0035] In one embodiment, the server device 150, or related device, has previously performed a crawl of the network 106 to locate articles, such as web pages, stored at other devices or systems coupled to the network 106, and indexed the articles in memory 162 or on another data storage device. It should be appreciated that other methods for indexing articles in lieu of or in combination with crawling may be used, such as manual submission.

[0036] It should be noted that other embodiments of the present invention may comprise systems having different architecture than that which is shown in Figure 1. For example, in some other embodiments of the present invention, the client device 102a is a stand-alone device and is not coupled to a network. The system 100 shown in Figure 1 is

merely exemplary, and is used to explain the exemplary methods shown in Figures 2 through 4.

[0037] Various methods in accordance with embodiments of the present invention may be carried out. For example, in one embodiment, an instant messenger event is captured by compiling event data associated with at least one instant messenger message, the instant messenger event is associated with a conversation, and at least some of the event data associated with the instant messenger event is indexed. In one embodiment, a search query is received, and the conversation is identified as relevant to the search query.

[0038] In one embodiment, capturing the instant message event can comprise identifying an activity associated with an instant messenger application on a client device, identifying the instant messenger event, and compiling the instant messenger event from at least some of the event data. The event data can comprise one or more of sender data, recipient data, a time associated with the event, a date associated with the event, and content from the instant messenger message associated with the event. Indexing at least some of the event data can comprise associating an event ID with the event and associating the event ID with at least some of the event data. The instant messenger event is captured on a client device or in another embodiment on a network device.

[0039] In one embodiment, the instant messenger event can be compiled upon the sending or receipt of an instant messenger message. The instant messenger event can be compiled after a period of time. The period of time can be predetermined or it can be a period of inactivity on the instant messenger application.

[0040] In one embodiment, identifying instant messenger activity can comprise one or more of identifying instant messenger network activity, identifying a user interface change associated with an instant messenger application, and determining that the instant messenger application is active. In one embodiment, identifying the instant messenger event can comprise monitoring the instant messenger application for an ongoing period of time or analyzing a current state of the instant messenger application to identify the instant messenger event. Identifying the instant messenger event can comprise one or more of identifying a display area associated with the instant messenger application and determining the content of the display area monitoring operating system calls made by the instant messenger application to display text, hooking into the instant messenger application's notification application program interface, directly querying the instant messenger application, and extracting text from a display area associated with the instant messenger application.

[0041] In one embodiment, associating an instant messenger event with a conversation can comprise determining if an existing conversation relevant to the instant messenger event exists, associating the instant messenger event with an existing conversation if the existing conversation is determined to be relevant to the instant

messenger event, and associating the instant messenger event with a new conversation if no existing conversation is determined to exist that is relevant to the instant messenger event. In one embodiment, determining if an existing conversation exists is based at least in part on participants in the message and a time the message was received or sent. In one embodiment, a title associated with the conversation can be determined.

[0042] Figure 2 illustrates an exemplary method 200 that provides a method for capturing and processing an event. This exemplary method is provided by way of example, as it will be appreciated from the foregoing description of exemplary embodiments there are a variety of ways to carry out methods in other embodiments of the present invention. The method 200 shown in Figure 2 can be executed or otherwise performed by any of various systems. The method 200 is described below as carried out by the system 100 shown in Figure 1 by way of example, and various elements of the system 100 are referenced in explaining the example method of Figure 2.

[0043] In 202, the capture processor 124 captures an event. The event can be a real-time event or can be a historical event. The capture processor 124 can capture a real-time event by identifying and compiling event data associated with the event upon the occurrence of the event. The capture processor 124 can capture a historical event, for example, by periodically crawling the memory 108 or associated data storage device of the client device 112a for previously uncaptured articles or receiving articles or data from client applications and identifying and compiling event data associated with the event. The capture processor 124 may have separate capture components for each client

application, network monitoring, performance data capture, keystroke capture, and display capture. In one embodiment, the capture component can use a generalized event definition mechanism, such as an event schema that it has previously defined and registered with the client device 102a, to capture or express the event.

[0044] Figure 3 illustrates an exemplary method 202 for capturing an instant messenger event. In 302, an instant messenger capture component identifies an instant messenger activity. An instant messenger activity can be any suitable activity that indicates that an instant message has been sent or received by an instant messenger application or that the instant messenger application is active on the client device. For example, an instant messenger activity can be instant message network activity, such as the receipt or sending of an instant messenger packet. In one embodiment, the instant messenger capture component can monitor instant messenger network activity, such as packets sent and received by the client device 112a to or from the network 106, and can identify instant messenger network activity. Instant messenger network activity, such as instant messenger packets, can be identified by a port on the client device associated with the packet(s) or by a header or signature contained in the packet(s). A single instant messenger message may be split across multiple packets. The identification of instant messenger network activity, such as an instant messenger packet, can indicate that the instant messenger application is in use. Alternatively, the capture components can examine the client device for network ports that are in listening mode to determine the presence of an active instant messenger application.

[0045] Other types of instant messenger activity can be used to determine that the instant messenger application is in use. For example, the instant messenger capture component can determine that the instant messenger application is active by examining the instant messenger application in memory, such as by detecting the presence of an application with an instant messenger class name, detecting the presence of an instant messenger display area output to the display of the client device, detecting that a display area associated with the instant messenger application is active or detecting any other suitable change in the client device user interface associated with the instant messenger application.

[0046] After the identification of instant messenger activity, in 304, the instant messenger capture component can attempt to identify an instant messenger event, such as for example the receipt of an instant messenger message, the sending of an instant messenger message, or entering text to compose an instant messenger message. The instant messenger capture component can identify an instant messenger event by analyzing or monitoring the instant messenger application. The instant messenger capture component can monitor the instant messenger application for an ongoing period of time to identify an instant messenger event or events or can analyze the current state of the instant messenger application to identify an instant messenger event or events.

[0047] In one embodiment, the instant messenger capture component can identify or locate a display area associated with the instant messenger application and determine the content of the display area, such as a display window. The instant messenger capture

component can, in one embodiment, monitor operating system calls made by the instant messenger application to display text to determine the content of the display area. In another embodiment, the instant messenger capture component can hook into the instant messenger application's notification application program interface (API) to determine the content of the display area. In another embodiment, the instant messenger capture component can directly query the instant messenger application for the content of the display area. The instant messenger capture component can also extract text from the display area to determine the content of the display area. In one embodiment, the instant messenger capture component can analyze visual aspects of the instant messenger application, such as a bitmap, to extract text from it using, for example, optical character recognition.

[0048] The instant messenger capture component can monitor logs generated by the instant messenger application and/or monitor keystrokes entered into the instant messenger application to determine the content of the instant messenger application. In the embodiment shown in Figure 1, for example, the instant messenger capture component resides on the client device and captures the instant messenger event on the client device. In another embodiment, the instant messenger capture component can be on a network device, such as a server, and the instant messenger event can be captured, indexed and/or stored on the network device.

[0049] In 306, the instant messenger capture component can capture the extracted event data associated with the instant messenger event. In one embodiment, the instant

messenger capture component can use an event schema to express the event data. For example, the event schema for an instant messenger event can include one or more of sender information, recipient information, time that the message was sent or received, a date that the message was sent or received, and the content of the message. The event schema can also contain a conversation ID. As explained above, a conversation ID can indicate the particular conversation that the instant message is associated with. A conversation can end with a user logging out or closing the instant messenger application or the instant messenger application being inactive for a period of time. The conversation ID can be provided by the instant messenger application. Alternatively, the conversation ID can be generated by the instant messenger capture component or the indexer 130. In this embodiment, the instant messenger capture component can determine if instant messenger messages are part of the same conversation based on some or all of the recipient(s) and sender of the messages, the time in between receipt or sending of messages, and whether messages are in the same instant messenger application display area. For example, if user 112a sends a message to a user B and then receives a message from user B two minutes later, the instant messenger capture component can associate the two messages with the same conversation and assign each the same or similar conversation ID.

[0050] The instant messenger capture component can compile an event upon the receipt or sending of an instant messenger message. Alternatively, the instant messenger capture component can compile an event after a predetermined amount of time (for

example, 2 minutes) from the receipt or sending of a message or from the last instant messenger event. In another embodiment, the instant messenger capture component can compile an event after a period of inactivity. The period of inactivity can vary based on the particular user or can be a predefined period of time.

[0051] Returning to Figure 2, in 204, the capture processor 124 determines whether the event captured is an indexable event. As explained above, some real-time events may not be indexed (non-indexable real-time events). In one embodiment, non-indexable real-time events are used to update the current user state and are, for example, examining a portion of an article, changing an article, and closing an article. In this embodiment, non-indexable events are not indexed or sent for storage by the indexer 130. Indexable events can be indexable real-time events or historical events.

[0052] If an indexable event is determined, then, in 206, the event can be sent by the capture processor 124 to the queue 126 with an indication that it is an indexable event. In the embodiment shown in Figure 1, indexable real-time events are sent to both a user state queue and an index queue within queue 126 and historical events are sent to the index queue within the queue 126. Alternatively, indexable real-time events may not be sent to the user state queue to save computational time. The capture processor 124 can send the event in a form described by an event schema to the queue 126. If the event is determined to be a non-indexable event, then, in 206, the non-indexable event can be sent by the capture processor 124 to the user state queue of the queue 126 with an indication that it is not to be indexed.

[0053] In one embodiment, the queue 126 holds the event until a condition is met, such as the search engine is ready to receive it. Based on the event data, the event can be prioritized on the queue 126 for handling. For example, historical events are given a lower priority for processing by the queue 126 than real-time events. In one embodiment, when the indexer 130 is ready to process another event, it can retrieve an event or events from the index queue in the queue 126. The query system 132 can retrieve an event or events from the user state queue of the queue 126, when it is ready to update the user state. In another embodiment, a queue is not used and events are sent directly to the search engine 122 from the capture processor 124.

[0054] In 208, the indexer 130 indexes and stores the event. The indexer 130 can retrieve an event from the queue 126 when it is ready to process the event. In one embodiment, the indexer 130 determines if the event is a duplicate event and if not assigns an Event ID to the event. The indexer 130 can also associate the event with related events. In the embodiment shown in Figure 2, the indexer determines indexable terms associated with the event, dates and times associated with the event, and other data associated with the event from the event schema. The indexer 130 can associate the Event ID with the indexable terms that are contained in the index 142. The event can be stored in the database 144 and the content of the event can be stored in the repository 146.

[0055] Figure 4 illustrates an exemplary method 210 for indexing and storing an instant messenger event. In 402, the indexer 130 retrieves an instant messenger event

from the queue 126 and determines that the event is an instant messenger event. In one embodiment, the indexer 130 can determine if the retrieved instant messenger event is a duplicate of a previously processed instant messenger event, and if not can assign the instant messenger event an event ID. The indexer 130 can also parse out indexable terms, including sender and recipient names, times and dates associated with the event, and a conversation ID from the event schema. In one embodiment, the indexer 130 determines a title associated with an instant messenger conversation. The indexer 130 can determine the title based on the frequency of unique terms appearing in the conversation, the first sentence of the conversation, or the participants and time of the conversation. For example, the indexer 130 can select the first sentence in the first message of the conversation and associate that sentence as the title of the conversation. In another embodiment, the query system or other component can do this in a post-processing phase.

[0056] In 404, the indexer 130 determines whether the instant messenger event is associated with a new conversation or if the instant messenger event is associated with an existing conversation. As explained above, the instant messenger event can have a conversation ID associated with it. The indexer 130 can compare the conversation ID with previously processed instant messenger event conversation IDs to determine if there is a match. If no conversation ID exists, the indexer 130 can determine if the instant messenger event is associated with an existing conversation based on some or all of the recipient(s) and sender of the message, the time in between receipt or sending of the

message and previous messages, content of the message and messages associated with a conversation, and whether the message is associated with the same instant messenger application display area as previous messages. For example, if the indexer 130 has previously processed an instant messenger event where user 112a sent a message to a user B and is now processing an instant messenger event where user 112a receives a message from user B two minutes later after the first event, the indexer 130 can associate the two events with the same conversation. Additionally, if user A sends user B a number of messages in a row, these messages may be associated with the same conversation. In other embodiment, messages between multiple users, such as in a “chat room” can be associated with the same conversation, messages from a user to himself can be associated with the same conversation, and messages from a user to an automated service or computer can be associated with the same conversation.

[0057] If the indexer 130 determines that the instant messenger event is associated with a new conversation, then the indexer, in 406, can create a new conversation object. A conversation object can associate an instant messenger event with related instant messenger events from the same conversation. Conversations can also be associated with each other through an associated conversation identifier.

[0058] If the indexer 130 determines that the instant messenger event is associated with an existing conversation, then the indexer 130, in 408, can load the relevant conversation object from the database 144 of the data store 140. By associating instant

messenger events with a conversation, a user is able to search for and retrieve complete instant messenger conversations.

[0059] Once the event has been associated with a conversation object and provided with a conversation object identifier, the instant messenger event can be stored in the data store 140. In 410, the instant messenger message is stored in the repository 146. In one embodiment, if the instant messenger application stores the conversation, a link or any suitable associated identifier for the location of the stored conversation can be stored in the data store associated with the Event ID. In 412, the event and conversation object are stored in the database 144. In 414, the indexable terms from the instant message are stored in the index 142. In one embodiment, the Event ID from the event is associated with terms in the index 142 that equate to the indexable terms of the event.

[0060] The capturing of instant messenger message events that are indexed and stored by the search engine 122 allows the user 112a to search for and retrieve instant messenger messages on the client device 102a and allows the search engine to automatically search for instant messenger messages on the client device 102a. Indexing instant messenger events while they happen without waiting for the end of the conversation provides the user with access to the event data before the end of the conversation.

[0061] The environment shown reflects a client-side search engine architecture embodiment. Other embodiments are possible, such as a stand-alone client device or a network search engine.

[0062] While the above description contains many specifics, these specifics should not be construed as limitations on the scope of the invention, but merely as exemplifications of the disclosed embodiments. Those skilled in the art will envision many other possible variations that are within the scope of the invention.